

Standards-Based Wireless Sensor Networking Protocols for Spaceflight Applications

Raymond S. Wagner, Ph.D.
NASA Johnson Space Center
2101 NASA Parkway, Mail Code EV814/ESCG
Houston, TX 77058
281-244-2428
raymond.s.wagner@nasa.gov

Abstract—Wireless sensor networks (WSNs) have the capacity to revolutionize data gathering in both spaceflight and terrestrial applications. WSNs provide a huge advantage over traditional, wired instrumentation since they do not require wiring trunks to connect sensors to a central hub. This allows for easy sensor installation in hard to reach locations, easy expansion of the number of sensors or sensing modalities, and reduction in both system cost and weight. While this technology offers unprecedented flexibility and adaptability, implementing it in practice is not without its difficulties. Recent advances in standards-based WSN protocols for industrial control applications have come a long way to solving many of the challenges facing practical WSN deployments. In this paper, we will overview two of the more promising candidates – WirelessHART from the HART Communication Foundation and ISA100.11a from the International Society of Automation – and present the architecture for a new standards-based sensor node for networking and applications research.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. CHALLENGES OF WIRELESS SENSING.....	2
3. OVERVIEW OF WSN STANDARDS.....	2
4. EXPERIMENTAL METHODOLOGY FOR COMPARING STANDARDS.....	5
5. CONCLUSIONS	7
REFERENCES	7
BIOGRAPHY	7

1. INTRODUCTION

Wireless Sensor Networks (WSNs) offer a new paradigm for acquiring sensor data. Rather than gathering sensor data through wired data buses, WSNs employ a wireless backhaul to transmit sensor readings to central locations for aggregation and further processing. This provides a number of potential benefits in spacecraft design, not the least of which is the potential to substantially reduce system weight by eliminating wiring harnesses and connectors. Un-tethering sensors from wires also opens up a new range of possibilities. Sensing infrastructure need no longer be fixed following spacecraft design and manufacture: should situational awareness be enhanced

by re-locating a sensor from one panel to another, this is easily accomplished. Similarly, additional sensors can be easily added to the existing suite, providing a more detailed measurement set without requiring more wires to be strung behind bulkheads and through walls of pressure shells. Finally, wireless sensors can be re-used between vehicles once their initial missions have been ended. A WSN node can be relocated from a spent vehicle, such as a lunar lander, to one currently in service, such as a lunar rover or habitat. The node can even be outfitted with a new set of sensors in the process, retaining the common radio and networking hardware, to give a new functional unit built mostly from recycled parts. Re-purposing wired systems would be much more difficult, requiring wiring to be stripped from one craft and re-strung in another, necessitating substantial disassembly of spacecraft in both cases.

While this technology offers unprecedented flexibility and adaptability, implementing it in practice is not without its difficulties, particularly with respect to achieving reliability that is on par with wired sensor approaches. Any practical WSN deployment must contend with a number of difficulties in its radio frequency (RF) environment including multi-path reflections and interference from other systems. Techniques must be designed to overcome all these factors, while at the same time operating at a low enough power draw to allow operational lifetimes on the order of years using small, onboard batteries.

In recent years, a great deal of focus has been given to solving these common problems for WSN applications in industrial automation and control, where the modern factory, refinery, or offshore drilling platform presents an incredibly challenging RF environment. These efforts by government, academic, and industrial partners have resulted in standards-based wireless sensor network (SB-WSN) protocols capable of communication reliability approaching that of wired solutions with a very low per-node power consumption and network lifetimes approaching the decade mark. Given similarities in operational requirements between mission-critical industrial processes and spaceflight applications – namely, the insistence that data transport be both extremely reliable as well as timely – we maintain that these SB-

WSN protocols hold great promise in the aerospace arena as well.

In this paper we will overview the two major standards to emerge from the industrial control field – WirelessHART from the HART Communication Foundation and ISA100.11a from the International Society of Automation (ISA). Both are rooted in the IEEE 802.15.4 standard and provide a level of robustness and reliability that should make them well suited to spaceflight applications. We will provide a technical review of both protocols starting with their inspiration as a means to extend the capabilities of ZigBee – the first commercial 802.15.4-derived protocol – which has had limited uptake in the mission-critical industrial control market. We will then discuss the methodology of an in-house performance evaluation of these protocols in a controlled environment and present preliminary results.

2. CHALLENGES OF WIRELESS SENSING

Because sensor nodes are designed to operate without wire interconnects for data transfer, they will typically not have wired power connections. This means nodes must rely on local power sources such as power scavenging or onboard batteries for both data processing and communication. Sensor nodes must often be small and have service lives on the order of years, so this necessitates very low power operation. With wireless communication representing the largest portion of a node's power utilization, the node must therefore restrict itself to very low-power transmission with periodic sleeping/waking of radio circuitry. Such low-power radio-frequency (RF) communication is extremely vulnerable to a variety of distortion and interference mechanisms. Multi-path reflections can distort signals, limit data rates, and cause signal fades that prevent nodes from having clear access to channels, especially in a closed environment such as a spacecraft. Other RF signal sources, such as wireless internet, voice, and data systems may contend with the sensor nodes for bandwidth. Finally, RF noise from electrical systems and periodic scattering from moving objects such as crew members can collectively contribute to a highly unpredictable, time-varying communication environment.

Communication reliability is key when replacing wired systems with wireless equivalents, so to cope with this difficult RF environment, a WSN must rely on several different mechanisms to deliver reliable performance. For example, it must utilize intelligent channel access mechanisms in order to constantly monitor and exploit channel quality and availability, which may vary rapidly with time. In addition, it must employ an intelligent routing mechanism that can move data reliably from a source node back to the central data manager, possibly utilizing multiple different paths relayed through multiple

different intermediate nodes. Finally, it must provide simple, reliable mechanisms to expand and contract (scale) the network, allowing nodes to enter and leave as necessary and changing the routing structure accordingly. Moreover, in addition to providing all these services, the WSN must be designed so that nodes make minimal use of their radios, both for data transfer and network coordination, to permit years-long operational lifetimes.

Designing such robustness at the outset is incredibly challenging, and it requires expertise at all layers of the networking stack, from physical radio design to channel access schemes, routing protocols, and distributed data-processing algorithms. Fortunately, a critical mass of effort across a variety of fields in wireless sensing has emerged in the last few years, resulting in the development of WSN standards that can be leveraged by spaceflight applications. We will now provide a brief overview of these standardized protocols.

3. OVERVIEW OF WSN STANDARDS

Robust, reliable wireless sensor/actuator networks stand to benefit a number of industries, and as a result much effort has been expended in recent years to develop design standards for WSNs addressing many of the aforementioned problems. The core of these efforts is the IEEE 802.15.4 standard, which targets low-data rate applications requiring wireless interconnections between measurement, analysis, and control devices – aggregations which can be classified as personal area networks (PANs). This puts 802.15.4 in the same general family as the Bluetooth standard (IEEE 802.15.1), though the low-power and low-data rate nature of its intended applications differentiates it from the latter. An 802.15.4 PAN can be analyzed using a simplified version of the Open System Interconnection (OSI) protocol stack consisting of the following layers (bottom to top): physical (PHY), medium access control (MAC), networking (NET), and application (APP). 802.15.4 specifies only the PHY and MAC layers. The remaining layers are provided by subsequent protocols such as ZigBee, WirelessHART, ISA100.11a.

The 802.15.4 PHY specification requires radios to operate in one of three frequency bands: 868-868.8 MHz (Europe, one channel), 902-928 MHz (North America, 30 channels), and 2400-2483.5 MHz (worldwide, 16 channels). A number of modulation schemes are allowed in the original 2003 standard and the 2006 and 2007 updates, but the most common are flavors of direct sequence spread spectrum. To regulate access to the channel, the 802.15.4 MAC describes a carrier sense multiple access with collision avoidance (CSMA-CA) scheme. That is, a device with a frame to send will first listen to the channel, and if there is no activity it can begin transmitting its data. If the sending device finds that the

channel is already in use, it waits for a random period and then checks the channel for activity again, either transmitting or waiting for another random number period depending on what it observes. Typically some limit on the number of attempts to make will be set so that devices do not wait ad infinitum to send a single frame [1-3].

Standardized WSN networking stacks build on top of the PHY and MAC provided by 802.15.4, adopting those two layers either outright or with some modification. We will now discuss the three most prominent alternatives, starting with ZigBee.

ZigBee

ZigBee is a protocol that is more-or-less designed to ride directly above the 802.15.4 PHY and MAC layers, providing NET and APP layers to yield a complete protocol stack. As is the case with all 802.15.4-based systems, ZigBee is designed for low-power, low-data rate applications. The ZigBee protocol has gone through several iterations to date, the most recent of which is the ZigBee 2007 specification. ZigBee 2007 defines two protocol stacks. The first is simply called “ZigBee” and is very similar to the previous single-stack ZigBee releases. It is designed for light-duty use in the home and the office (e.g., home lighting control). The second, called “ZigBee PRO”, is a more robust protocol designed for industrial control applications. It provides more reliable performance but requires implementation of a larger and more complicated protocol stack [4].

ZigBee defines three classes of devices: ZigBee Coordinators (ZC), ZigBee Routers (ZR), and ZigBee End Devices (ZED). Each network has one ZC, which is responsible for network formation and which can also aid in message routing. ZR’s also participate in routing and can run a sensing/actuation application as well. ZED’s only run applications and cannot participate in message routing – each ZED must report to either a ZR or the ZC [5,6].

ZigBee uses the 802.15.4 PHY and MAC layers directly, though the 2007 version does allow for some limited frequency agility in the PHY layer, so that radios can be automatically switched away from problem channels when throughput falls off. Regarding the NET layer, both ZigBee Coordinators and ZigBee Routers participate in multi-hop routing of messages; ZigBee End Devices only address messages to their associated parent routing device, which is found within their radio transmission range. With these device roles, a ZigBee network can have one of three topologies: (1) star, where all non-coordinator devices report directly to the coordinator; (2) tree, where all ZOD’s report to a routing device, and routing devices communicate up and down a tree of routing devices in a well-defined hierarchy crowned by the ZC; and (3) mesh, where both ZOD’s and routing

devices are free to communicate with any other routing device within radio range. While ZODs are allowed to periodically cycle into a low-power sleep mode in networks using 802.15.4’s limited duty cycling feature, ZR’s must in general always remain awake. Some limited provision for ZR duty cycling does exist using the 802.15.4 “beaconing” mechanism, which attempts to establish a crude synchronization among nodes, though the period of sleep/wake cycles is limited by the inability of low-cost ZigBee hardware to maintain precise timing. Beaconing in general also requires longer wake times from ZEDs [4-6].

While ZigBee has found a market in home and office settings, the protocol has not been as widely embraced by the industrial process measurement and control industry, even in its more robust ZigBee PRO form. It has been found that the solely contention-based MAC is not able to reliably provide the message delivery required by critical industrial applications. Since 802.15.4-based sensor nodes occupy the shared industrial, scientific, and medical (ISM) RF band, they can expect to observe transmissions for a variety of protocols with higher radiated power, such as IEEE 802.11 and Bluetooth. In the presence of such traffic, the 802.15.4 MAC will always back off, potentially leaving the nodes unable to get the channel access required to send time-critical messages in a timely manner. This is exacerbated in situations where a great deal of multi-path reflection is to be expected [10,11].

As a result, a pair of standards for high reliability wireless networking of sensors in very difficult RF environments has emerged.

WirelessHART

The first, WirelessHART, is specified in the HART Field Communications Protocol, a long-established device measurement and control protocol in industrial automation. Revision 7 of the protocol augments the formerly wired-bound HART with a wireless data delivery mechanism based on the 802.15.4 physical layer. WirelessHART is designed from the ground up to enable wireless sensing and actuation in very harsh industrial environments, where communication over the wireless network must have reliability comparable to wireline communication. It also assumes that wireless assets will be deployed where wired assets are difficult to place, so that WirelessHART devices must be able to operate for a long time on a single set of batteries.

WirelessHART uses the 802.15.4 PHY as-is. It operates in the 2.4 GHz band and employs DSSS channel coding. A significant departure is taken from the 802.15.4 MAC in the WirelessHART specification, however. The WirelessHART MAC, based on the Time-Synchronized Mesh Protocol (TSMP) originally developed by Dust

Networks, Inc., employs time-division multiplexing of the channel rather than the carrier-sensing and random backing off of the 802.15.4 MAC. This alternative MAC design was motivated by the extremely hostile radio-frequency (RF) environments that WSN's are likely to encounter in industrial environments. Industrial deployments can be expected to be plagued by RF interference from other wireless systems such as Wi-Fi networks and cordless phones, RF noise from machinery, physical obstruction of radio paths between devices, multipath effects between sources and receivers, and node losses due to depleted battery supplies and environmentally unfriendly operating conditions. In addition, these effects are likely to be highly time-variant, precluding an approach that attempts to compensate by carefully calibrating a network to account for conditions at the time of its deployment. A robust, agile system capable of working around changing ambient conditions must be designed [7].

The cornerstone of the WirelessHART MAC is time division multiple access (TDMA) to the channel, rather than the CSMA-CA approach taken by the 802.15.4 MAC. This first requires time synchronization among nodes, which is maintained by embedding time offset information in acknowledgement (ACK) packets sent to confirm successful reception of messages. Piggybacking this service on ACK packets allows TSMP to avoid expensive beaconing approaches for synchronization such as in the 802.15.4 MAC. Once a pair of nodes is synchronized, a schedule can be established for communication. Time is divided into slots, and proper synchronization allows for agreement between the pair on slot start time [9].

Next, the nodes must decide which sub-band of the 16 available in the 2.4 GHz PHY they will use. To do so, they agree on a start point in a sequence of channels. Beginning at this agreed-upon channel, at each new slot the pair switches to the next channel in the sequence. If two nodes find that they are unable to communicate on a poor-quality channel in a given slot, they need only to wait one slot-length for a new chance to communicate. Blacklisting of channels with repeatedly poor performance is supported so that these can be skipped in the slotted sequencing of channels [9].

Each receiving node in TSMP sends an ACK to the sender upon successful reception. Should the sender not receive an ACK, it may switch to an alternate neighboring node on its "parent" list and re-try the transmission. This parent list records valid next-hop neighbors for a given destination and is formed when the sending node initially joins the network. Each parent's entry in the sending node's parent list is mirrored by the node's entry in that parent's corresponding "child" list. Since alternate parents occupy different points in space, this adds an element of spatial diversity to the overall MAC scheme.

Furthermore, since the re-try to the alternate parent will be at a later time-slot and on a different frequency, the scheme also features time diversity and frequency diversity. Finally, the direct sequence spread spectrum channel coding in the PHY layer adds code diversity, for a system with a great deal of diversity and, hence, agility in coping with difficult wireless environments [9]. In fact, data transport in a well-formed WirelessHART network is typically greater than 3-sigma (99.7300204%) reliable, and under normal circumstances is greater than 6-sigma (99.9999998%) reliable [8].

All nodes in a WirelessHART network are full-function devices capable of routing multi-hop traffic. Thus, all WirelessHART networks form full mesh network topologies. Note that, while each node is a router, each can also be efficiently battery powered due to the aggressive duty-cycling that the TSMP protocol allows (< 1% time full-power) [7]. Even routers can periodically be put into a low-power "sleep" state with its radio powered down, since time synchronization guarantees that when nodes "wake" and power up their radios, they will all be doing it at the same time. Graph routing tables are computed by a central network manager and distributed to nodes on routing paths so that they know which neighbors are their next-hop parents on the path to the requested destination node. The number of parent choices available at each node can be varied depending on the criticality of the route [9].

Finally, WirelessHART supports application messages which conform to the HART device communication protocol. In an adaptation of WirelessHART to spaceflight applications, however, this layer could be ignored, instead treating the payloads of WirelessHART packets as generic buffers for application data.

ISA100.11a

WirelessHART has become the first standard to the market in industrial automation, but following its introduction an effort began in ISA to develop a standard suitable to a wider class of plant control networks than just those made of HART devices. The first standard to come out of the ISA100 wireless group is ISA100.11a, which provides a wireless backend suitable to use with all manner of legacy device communication protocols at the application layer. ISA100.11a is in the final stages of ratification within ISA100 and should be released sometime in late 2009 or early 2010.

ISA100.11a shares many aspects in common with WirelessHART, including a TDMA MAC scheme based on Dust Networks' TSMP algorithm, and in fact many features found in WirelessHART are designed into ISA100.11a. The ISA protocol aims to provide a larger set of options, however, for industrial WSNs. Specifically, ISA100.11a re-introduces a contention-based

MAC along the lines of original 802.15.4/ZigBee specification in addition to the WirelessHART-like TDMA-based MAC to enable higher throughput when desired. While this CSMA-CA option will suffer the same ill effects as ZigBee in difficult RF environments, it will allow sensor nodes to have greater bandwidth in more friendly environments. Thus, balancing of reliability and throughput is possible, with the option of trading reliability for throughput in more harsh environments or achieving higher throughput in environments that are more friendly to low-power wireless networking. Additionally, ISA100.11a supports a more simple class of non-routing device, whereas WirelessHART requires that all devices be capable of routing network traffic. While this limits the number of alternative paths available in the network, it does allow designers to trade off device cost for routing redundancy. Several other distinctions of varying importance exist between the two protocols, such as details regarding security, wired plant backbone networks, and inclusion of packet fragmentation and re-assembly (included in ISA100.11a but not WirelessHART). [12]

The co-existence of these two advanced industrial WSN protocols seems likely. A working group within ISA100 has been investigating a merger of the standards, and this continues to remain an eventual possibility. In the near term, dual-boot devices are likely to provide an interim solution to co-existence: nodes may be able to boot either the ISA100.11a stack or the WirelessHART stack as designated during their commissioning, with code blocks common to the two protocols shared by the implementations of each.

4. EXPERIMENTAL METHODOLOGY FOR COMPARING STANDARDS

We began our investigation of IEEE 802-15.4-based protocols by using off-the-shelf WSN hardware from Crossbow, Inc. These parts used the 802.15.4 PHY and MAC layers directly and were configured in a star topology with each node reporting directly to the network base station. This configuration is analogous to an instance of the ZigBee protocol stack, configured in a star topology, below the application layer. Moreover, nodes were operated at their full duty cycle, so that they did not periodically enter a low-power “sleep” state. Nodes were mounted in the Lunar Habitat Wireless Testbed (LHWT) shown in the foreground of Figure 1. The LHWT resides at Johnson Space Center in Houston, TX and serves as an environment for testing the co-existence of multiple wireless systems in a closed, reflective environment similar to that expected in a habitable environment on the lunar surface or on orbit.

For, a test application, we chose wireless micrometeoroid orbital debris (MMOD) impact detection, which provides quite a challenging problem for a wireless data acquisition

system. Nodes monitor single-axis accelerometers mounted perpendicular to the habitat hull for high frequency, transient signals corresponding to impact events. Signals may then be analyzed for features of interest useful for solving problems such as impact localization. In this example, however, nodes merely attempted to acquire sample points from the transient impact signals at the maximum sample rate afforded by their onboard analog-to-digital (A/D) converters. Measurements were summarized into packets that were then sent to the base station for display and logging.

The results were rather disappointing, and stemmed from two main sources: (1) lack of communication reliability afforded by the 802.15.4 MAC and (2) the processing architecture of the sensor nodes used in the experiment. Regarding the communication reliability, when the network was configured to operate in a 2.4 GHz sub-channel that happened to be shared with another wireless system (e.g., an 802.11 local area network), nodes were rarely able to send packets to the base station. In many cases, nodes dropped connectivity with the base station entirely. Even when the network was re-configured to broadcast in a less crowded 2.4 GHz sub-band, nodes would occasionally suffer excessively long communication latencies with the base station. This behavior confirms the criticism of the 802.15.4/ZigBee MAC with regards to the lack of frequency agility in the available sub-bands.

Regarding the processing architecture, the high-bandwidth requirements of MMOD data acquisition proved to be quite illuminating. We used a calibrated “pinger” to generate impacts on the hull, and we verified through the use of accelerometers and a wired data acquisition system that the impacts generated highly repeatable accelerometer traces across multiple experiments. However, when analyzing the results gathered from the wireless sensors, we found that the recorded signals were far from uniform. In several cases, nodes missed the transient event entirely, reporting only background signals over the course of the experiment. In others, samples indicated widely varying peak magnitudes, although maximum values should be nearly consistent when using the calibrated pinger.

Although the accelerometers used with the Crossbow nodes do not have a high enough bandwidth to capture transient impact signals without some under-sampling, aliasing is not sufficient to explain the cases where nodes missed ping events entirely. Further investigation into the architecture of the code running on the nodes revealed the likely cause to be the method of task scheduling on a node's embedded microprocessors. Simply put, to capture high-frequency events such as MMOD impacts, nodes must be able to sample their sensors constantly without interference from other tasks, such as servicing the network stack. When nodes use a single microprocessor



Figure 1: Lunar Habitat Wireless Testbed (foreground) at Johnson Space Center in Houston, TX.

to co-ordinate both data acquisition and networking, high-frequency transient events are likely to be missed either in part or in total. This motivates a new sensor node architecture more suited to both robust signal acquisition and robust networking.

Forward Plan

In response to these shortcomings, we have designed a new wireless sensor platform which should serve as a robust research and development tool for evaluating WSNs in a variety of contexts. The platform is highly modular, allowing components to be changed out to enable investigations concerning competing WSN protocol stacks, data processing algorithms, and sensor modalities. The design is shown in the diagram in Figure 2.

The core of the new node, referred to in the figure as the main board, manages the data acquisition and processing. It contains a low-power microcontroller, off-chip expandable memory, and the node's power supply. The microcontroller on the main board schedules sampling of the board's sensors (using its own A/D module) and processes the sensed data prior to transmission.

The main board interfaces with a radio module. Typically a commercial off-the-shelf (COTS) component, the radio module consists of a second microcontroller, radio circuitry, and antenna assembly. The radio module fully implements the networking stack for the protocol under investigation and interacts with the main board mainly through transmit and receive commands over a hardware interface (e.g., serial bus).

With an integrated main board and radio module, the core of a multi-purpose sensor node is complete. We can then add an application-specific sensor card, containing all the sensors needed from the application under investigation as well as auxiliary hardware such as digital signal

processing chips or more advanced A/D components when application requirements exceed the computation and signal processing capabilities of the main board's microcontroller. In such a scenario, the microcontroller may be used primarily to schedule the distributed sensing tasks, leaving the bulk of data acquisition and analysis to these more powerful secondary modules.

The goal of this modular design is twofold. First, it allows us to build a WSN "development kit" that can quickly be customized to meet new distributed sensing needs for applications research. Second, it allows us to interchange radio modules implementing different WSN protocol stacks while running the same front-end sensing applications to meaningfully compare the capabilities of different protocols and standard implementations.

We plan to interface this design with radio modules implementing both the WirelessHART and ISA100.11a stacks, for which COTS components are now coming to market. Using this platform, we intend to pursue the following lines of investigation:

- **RF issues:** How reliable is the data delivery of these advanced industrial WSN protocols in practice? How resistant are transmissions to multi-path and other RF interference? What are achievable throughput rates? How well does this system co-exist with other 2.4 GHz devices such as wireless LAN?
- **Power issues:** How power-hungry are these protocols in practice – do they achieve operational lifetimes in excess of five years as advertised? How does the sensing task affect battery lifetimes? How do scheduled and event-driven sensing differ in their power requirements?
- **Application issues:** How feasible is it to accurately sense high-frequency transient events? Can timing information derived from the CDMA MAC be used to accurately synchronize time-stamping of measurements across nodes?
- **Protocol issues:** How will future protocols improve upon WirelessHART and ISA100.11a (e.g., next-generation ZigBee)? When the market matures, which standard or standards will be best suited to spaceflight applications? Will modifications of standards be necessary?




Figure 2: architecture of sensor node under development; radio and sensing modules interface with a main controller board.

5. CONCLUSIONS

Standards-based WSN protocols developed for demanding industrial environments hold much promise for aerospace applications. They have been designed to cope with extremely difficult RF environments to provide highly reliable data delivery and operational lifetimes on the order of years using only onboard batteries. Preliminary research verifies the shortcomings of protocols relying solely on the CSMA-CA MAC of 802.15.4, such as ZigBee, and encourages further investigation into newer protocols with TDMA MAC options, such as WirelessHART and ISA100.11a. Using a modular platform dividing the processing load of sensing and networking between a pair of microcontrollers, we intend to fully investigate these protocols from both network and applications research angles.

REFERENCES

- [1] IEEE standard 802.15.4-2003, <http://standards.ieee.org/getieee802/download/802.15.4-2003.pdf>.
- [2] IEEE standard 802.15.4-2006, <http://standards.ieee.org/getieee802/download/802.15.4-2006.pdf>.
- [3] IEEE standard 802.15.4a-2007, <http://standards.ieee.org/getieee802/download/802.15.4a-2007.pdf>.
- [4] B.M. Blum, "ZigBee and ZigBee PRO: Which Feature Set is Right for You?", Embedded.com, Oct. 2008, <http://www.embedded.com/design/210700316>.
- [5] J.K. Young, "Clearing Up the Mesh About Wireless Networking Topologies", Embedded.com, Aug. 2008, [http://www.embedded.com/design/networking/210200649\(part1\)](http://www.embedded.com/design/networking/210200649(part1)), [http://www.embedded.com/design/networking/210200676\(part2\)](http://www.embedded.com/design/networking/210200676(part2)).

- [6] "ZigBee", Wikipedia, <http://en.wikipedia.org/wiki/Zigbee>.
- [7] DUST Networks, "Technical Overview of Time Synchronized Mesh Protocol (TSMP)",
- [8] "WirelessHART Technical Data Sheet", HART Communication Foundation, http://www.hartcomm.org/protocol/training/resources/wiHART_resources/wirelesshart_datasheet.pdf.
- [9] "IEC 62591: Industrial communication networks – WirelessHart Communication Network and Communication Profile", International Electrotechnical Commission (IEC) draft document 65C/532/CD, 2009.
- [10] J. Song, S. Han, A.M. Mok, D. Chen, M. Lucas, M. Nixon, W. Pratt, "WirelessHART: Applying Wireless Technology in Real-Time Industrial Process Control", Proc. IEEE Real-Time and Embedded Technology and Applications Symposium, 2008, pp. 377-386.
- [11] T. Lennvall, S. Svensson, F. Hekland, "A Comparison of WirelessHART and ZigBee for Industrial Applications", Proc. IEEE Intl. Workshop on Factory Communication Systems, 2008, pp. 85-88.
- [12] "Wireless Systems for Industrial Automation: Process Control and Related Applications", International Society of Automation (ISA) Standard 100.11a, Draft 2a, 2009.

BIOGRAPHY

Raymond Wagner leads the wireless sensor network research and development program at NASA's Johnson Space Center, and he is involved in related programs for development of wireless communications systems for habitat and surface operations. He has been with Johnson Space Center since the Fall of 2008, prior to which he earned his Ph.D. in electrical engineering at Rice University in Houston, Texas, with a thesis concerning distributed data processing algorithms for wireless sensor networks. His research interests include wireless sensor networks, digital signal processing, computer networking, and wireless communications.